



Full length article

Predicting chemical hazard across taxa through machine learning

Jimeng Wu^{a,b,1}, Simone D'Ambrosi^{c,1}, Lorenz Ammann^d, Julita Stadnicka-Michalak^a,
Kristin Schirmer^{a,e}, Marco Baity-Jesi^{a,*}

^a Eawag, Überlandstrasse 133, CH-8600 Dübendorf, Switzerland

^b Department of Environmental Engineering, ETHZ, Zurich, Switzerland

^c Department of Statistics, Sapienza University of Rome, Piazzale Aldo Moro, 5, 00185 Rome, RM, Italy

^d Swiss Federal Institute for Forest, Snow, and Landscape Research WSL, Zürcherstrasse 111, CH-8903 Birmensdorf, Switzerland

^e School of Architecture, Civil and Environmental Engineering, EPFL, Lausanne, Switzerland



ARTICLE INFO

Handling Editor: Adrian Covaci

Keywords:

Machine learning

Acute toxicity

Ecotoxicology

Animal testing

In vivo testing

RASAR

Fish

ABSTRACT

We applied machine learning methods to predict chemical hazards focusing on fish acute toxicity across taxa. We analyzed the relevance of taxonomy and experimental setup, showing that taking them into account can lead to considerable improvements in the classification performance. We quantified the gain obtained through the introduction of taxonomic and experimental information, compared to classification based on chemical information alone. We used our approach with standard machine learning models (K-nearest neighbors, random forests and deep neural networks), as well as the recently proposed Read-Across Structure Activity Relationship (RASAR) models, which were very successful in predicting chemical hazards to mammals based on chemical similarity. We were able to obtain accuracies of over 93% on datasets where, due to noise in the data, the maximum achievable accuracy was expected to be below 96%. The best performances were obtained by random forests and RASAR models. We analyzed metrics to compare our results with animal test reproducibility, and despite most of our models “outperform animal test reproducibility” as measured through recently proposed metrics, we showed that the comparison between machine learning performance and animal test reproducibility should be addressed with particular care. While we focused on fish mortality, our approach, provided that the right data is available, is valid for any combination of chemicals, effects and taxa.

1. Introduction

One of the pillars of modern civilization is the ability to synthesize and/or use an enormous range of chemicals, which allow for new and improved products and serve as pharmaceuticals, pesticides, food additives and the like. However, their benefits need to be weighed against their risk. For that purpose, governments put risk assessment procedures in place decades ago, with the aim of evaluating the impact of chemicals on human and environmental health. In these assessment procedures, exposure, *i.e.* the level of chemical to which organisms are subjected, and hazard, *i.e.* the ability of chemicals to cause negative effects on organism/population health, are considered together as determinants of risk ([SETAC], 2018). Hazard assessment traditionally depends, to a large part, on animal tests using vertebrates, such as tests with mice, rats and fish, depending if assessments are for human or environmental health. These animal tests are ethically controversial. They are also

highly resource consuming (time, personnel, test material needed) and can, in fact, not comply with the demand of rapid testing of an ever increasing chemical universe in need of evaluation (Rovida and Hartung, 2009). On these grounds, there are global efforts to refine hazard assessments using alternative or supplementary strategies (Hamm et al., 2017; Sullivan et al., 2021). One of them is the employment of machine learning (A.H. Vo et al., 2019).

Machine learning (ML) is a term broadly applied to the use of computational algorithms to infer emerging patterns from data. In the field of chemical hazard assessment, one can mine chemical structural and physico-chemical information, along with effect data reported from animal toxicity tests. This comes often with the advantage that the domain of applicability is unrestricted, provided that appropriate training data is available.

One of the first applications of ML to hazard assessment was the prediction of acute oral toxicity in rats (LD50 values), using local lazy

* Corresponding author.

E-mail addresses: jimeng.wu@eawag.ch (J. Wu), kristin.schirmer@eawag.ch (K. Schirmer), marco.baityjesi@eawag.ch (M. Baity-Jesi).

¹ Equal contribution.

learning, achieving correlation coefficients of $R^2 = 0.72$ on a test set of 2896 compounds (Lu et al., 2014). Further ML approaches to predict rat acute oral toxicity by other investigators (Wu and Wang, 2018) improved the R^2 to 0.86, and using discrete outputs to measure accuracies, values reached close to 95% (Xu et al., 2017). However, with their focus on single species data, these studies do not cover the dimensionality of taxa and species sensitivity differences to be considered if one wants to know chemical impacts on another species, such as humans, or on the many species constituting an ecosystem (Furuhama et al., 2015; Basant et al., 2016; Ai et al., 2019), the latter of which is the focus of ecotoxicology. In this case, the use of ML is even more crucial, since in practice it is unthinkable to test all chemicals on all the species present in an ecosystem. Acute toxicity predictions have also been the focus of few applications of ML in ecotoxicology, though a critical lack of literature on machine learning in this area has been denounced (Miller et al., 2018). Li et al. aimed to predict toxicity categories for acute toxicity to fish of pesticides by making use of nine chemical descriptors (Li et al., 2017). Their dataset was rather small (~ 1000 entries), and dominated by two species, rainbow trout (RT) and bluegill sunfish (BS). They trained models on the RT and BS data separately, and also trained models by mixing all the available taxa together. In binary classification, their balanced accuracy (average between sensitivity and specificity, see Section 2.4) was lowest when excluding RT and BS, it was about 0.815 when restricting to a single taxon, and 0.825 when treating the whole dataset. However, it is difficult to state whether these differences are statistically significant, and, if so, whether it should be attributed to the smaller size of the dataset when training without the RT and BS. Tuulaikhuu et al. studied the EcoTox database, which contained enough data to allow for using taxonomy as a training feature (Tuulaikhuu et al., 2017). They trained random forests to predict acute fish toxicity data (LC50), obtaining an $R^2 = 0.85$. They identified fish species and the octanol–water partition coefficient ($\log P$ or $\log K_{ow}$) as the most important factors to explain species sensitivities differences. Despite the lack of a hold-out test set (in addition to train and validation), this result seems to suggest that in ecotoxicology we should be training machine learning models that infer how toxicity varies across taxa.

Indeed, when measuring the performances of different approaches, one of the fundamental questions is whether the measured differences are statistically significant. This issue was tackled by Hou et al. (2020), by taking an average of 10 independent hold-out validation sets. This procedure allowed an estimation of statistical fluctuations, but it also imposed a burden on the training data availability, since these hold-out test sets cannot be used for training.

Another fundamental question is how good should we expect machine learning models to be, since *in vivo* toxicity data is often noisy (*i.e.* repetitions some of the same experiment can provide very different results), and testing on noisy data cannot return a perfect performance. This issue was inspected by Luechtefeld et al. (2018), who, focusing on specific OECD² chemical testing guidelines (all concerned with rabbits, guinea pigs, mice, rats or Chinese hamster), adopted a metric to compare the performance of their machine learning model with the consistency of the experiments themselves. Through this metric, they concluded that their model, called Read-Across Structure Activity Relationship (RASAR), is "outperforming animal test reproducibility". RASAR models were developed for binary classification, and adopt a succession of unsupervised and supervised learning, which allows the exploitation of information on endpoints that are different from the target one. For example, one could predict the LC50 of a chemical by integrating a Mortality with a Behavior dataset. Given the great success of RASAR models, a natural question is whether they perform equally good in the datasets used in ecotoxicology, and how they compare to standard ML models.

Here, we perform an extensive analysis of ML models applied to an ecotoxicological dataset focused on fish acute toxicity data, with an eye on the extra information that can be obtained from the knowledge about the toxicity of the same compound on more than one species. We compare several standard machine learning models with the two variants of the RASAR model developed by Luechtefeld et al. (2018). We extend RASAR models to an arbitrary number of classes, which allows us to study both binary (more toxic versus less toxic) and 5-class (non-toxic, slightly toxic, moderately toxic, highly toxic, very highly toxic) classification. We carefully analyze the impact of introducing taxonomic and experimental information as extra input features in addition to the chemical, since this introduction allows for models with a wide taxonomic range of applicability, which can learn the patterns relating to toxicity and taxonomic/experimental variability. Finally, we compare our results with two indicators of the reproducibility of the experiments. These suggest that, given the level of noise in *in vivo* data (*e.g.* the same experiment repeated twice may have very different outcomes), we are close to the maximum achievable performances. However, this should not be taken as an indication that these models are outperforming animal test reproducibility.

2. Materials and Methods

2.1. Problem Setup

We are interested in predicting whether an experiment will result in a certain chemical being more or less acutely toxic on a given species. The outcome of the experiment is the LC50, which is the concentration of a chemical required to kill half a population, in our case 50% of fish in a 24 h, 48 h, 72 h or 96 h acute exposure scenario. We treat the problem as binary classification according to the European regulation (EC) No 1272/2008 on the classification, labelling and packaging of substances and mixtures (CLP Regulation) for acute toxicity data (<https://reach-online.eu/clp/en/annex-i-4-4.1-4.1.2.html>). To do so, we define the binary target label y_2 :

$$y_2 = \begin{cases} 1 & \text{if } LC50 \leq 1 \text{ mg}/\ell, \\ 0 & \text{if } LC50 > 1 \text{ mg}/\ell. \end{cases} \quad (1)$$

Our main goal is to deduce y_2 from the data, which we denote as \vec{x} . In F we show analogous results for five-class labels.

2.2. Data

The data, \vec{x} , contains information on chemicals, \vec{x}_{ch} , on the taxonomy of the tested organism, \vec{x}_{tax} , and on the experimental conditions, \vec{x}_{ex} :

$$\vec{x} = (\vec{x}_{ch}, \vec{x}_{tax}, \vec{x}_{ex}). \quad (2)$$

The vectors \vec{x}_{ch} , \vec{x}_{tax} and \vec{x}_{ex} have respectively 893, 5 and 6 components. To build x and the related set of labels [Eq. (1)], we take *in vivo* experiments on fishes from the Ecotox database (<https://cfpub.epa.gov/ecotox/>), provided by the Environmental Protection Agency of the USA. This database contains \vec{x}_{tax} and \vec{x}_{ex} , but it only contains the CAS identifier for the used chemical. In A we describe in detail the sources of the data, and which features are contained in \vec{x}_{ch} , \vec{x}_{tax} , and \vec{x}_{ex} .

The used information on the chemical has two different natures. We use some standard properties of each chemical (*e.g.* $\log P$, water solubility, etc.); and the Pubchem2D vectors, that identify chemicals through a one-hot encoded 881-dimensional vector (see A). In other words, we can write \vec{x}_{ch} as the union of the properties \vec{x}_{pr} and the Pubchem2D data \vec{x}_{pubc} ,

$$\vec{x}_{ch} = (\vec{x}_{pr}, \vec{x}_{pubc}). \quad (3)$$

² Organisation for Economic Co-operation and Development.

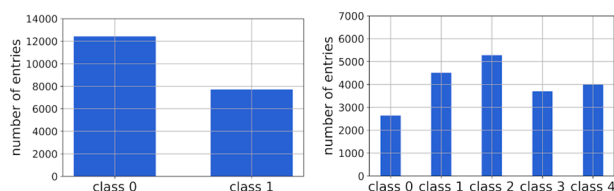


Fig. 1. Number of experiments per label after binary and 5-class labeling of the concentrations. The thresholds that define the classes are summarized in the appendix (Table 8).

We study *in vivo* mortality experiments performed on fishes, and take into account all the chemicals (mainly organic compounds) and species available in the dataset. As for the taxonomic ranks, we use class, order, family, genus and species as features. When the same experiment is repeated more than once, we merge the multiple entries into a single one, and take the median LC50 as an outcome. The data cleaning is described in A.

At the end of the process, we have $n_{\text{chem}} = 2199$ chemicals, tested for mortality on $n_{\text{spec}} = 345$ different species of fish. In total, our dataset counts 20128 entries. In Fig. 1, we show how the data is distributed into each class, in the case that we split the outputs into 2 (left) or 5 classes (right).

In addition to mortality, we also use data related to different effects, which is necessary as an input for the Data-Fusion RASAR model (B). This adds 3713 further entries to our dataset, spanning over 543 chemicals. Of those, 117 do not appear in the dataset related to mortality. See A for details.

2.3. Models

We train Logistic Regressions (LR), Random Forests (RF), K-Nearest-Neighbors (KNN), Multi-layer Perceptrons (MLP, this is a standard deep neural network model) (Alpaydin, 2020; Goodfellow et al., 2016), Simple RASAR (S-RASAR) and Data Fusion RASAR (DF-RASAR) (Luechtefeld et al., 2018). Details on the models are provided in B.

2.3.1. Euclidean Distances

The KNN and RASAR algorithms require measuring a distance between feature vectors. Initially, we will use Euclidean distances. In other words, the distance between two feature vectors \vec{u} and \vec{v} is

$$d(\vec{u}, \vec{v}) = \sqrt{(\vec{u} - \vec{v})^2}, \quad (4)$$

regardless of whether these vectors have ordinal or one-hot components.

2.3.2. Per-category Distances

The fact that our data is clearly separated into the three sources of information, x_{ch} , x_{ex} and x_{tax} , allows us to tailor our models to achieve better performance. One way to do this, is to define the distances in the KNN and RASAR models as a linear combination of the distances in different subspaces of the data.

We decide to treat separately the pubchem vectors (binary features), the chemical fingerprints (ordinal features), and the union of experimental and taxonomic details (these are categorical features, so we denote them with $\vec{x}_{\text{cat}} = (\vec{x}_{\text{tax}}, \vec{x}_{\text{ex}})$).

We measure the distance between pubchem vectors, through the Hamming distance function, d_{Ham} , which measures the number of equal entries in the two vectors. The distance between chemical fingerprints is a Euclidean distance, d_{Euc} , and for the categorical features we also use the Hamming distance. In order to have distances normalized to one (so that each set of features has the same weight), we divide the Hamming distances between two vectors by their number of components, and the Euclidean distances by their maximum over the training set. We call

\tilde{d}_{Ham} and \tilde{d}_{Euc} the resulting normalized distances.

The distance between two experiments \vec{u} and \vec{v} can then be written as

$$d(\vec{u}, \vec{v}) = \alpha_{\text{pr}} \tilde{d}_{\text{Euc}}(\vec{u}_{\text{pr}}, \vec{v}_{\text{pr}}) + \alpha_{\text{pubc}} \tilde{d}_{\text{Ham}}(\vec{u}_{\text{pubc}}, \vec{v}_{\text{pubc}}) + \alpha_{\text{cat}} \tilde{d}_{\text{Ham}}(\vec{u}_{\text{cat}}, \vec{v}_{\text{cat}}), \quad (5)$$

where α_{pr} , α_{pubc} and α_{cat} are constants whose optimal value is found through cross validation (CV).

Since we are interested in the relative (and not absolute) values of these constants, we can safely set $\alpha_{\text{pr}} = 1$. Then, to find the optimal values of α_{pubc} and α_{cat} , we perform a search on a grid in the plane (α_{pubc} , α_{cat}), tuning one dimension at a time, and choosing the alpha with the best accuracy. When two or more values of α give the same validation accuracy (with a tolerance of 0.001), we choose the smallest value. For both α_{pubc} and α_{cat} , we first found the optimum a coarse logarithmic grid, and then fine-tuned it with a finer one. Since we have several models that employ these distances, and because this hyperparameter tuning is computationally expensive, we only find the optimal α s for the K-NN algorithm, and use those values for the S and DF RASARs.

The advantage of using the distances in Eq. (5) is twofold:

1. It allows to assign a different weight to different subsets of features. This is useful because we do not expect all the groups of features to be equally important. Splitting the features into smaller groups can further improve the performances, but this would be at the expense of a longer hyperparameter tuning.
2. Since the distances d_{Euc} , d_{pubc} and d_{cat} are normalized, the constants α automatically tell us the relative importance of each source of data. In particular, the value of α_{cat} gives us a direct indication of the role played by taxonomy and experiment details.

2.4. Training and Assessing Performance

We divide our dataset in a training and a hold-out test set with an 80:20 ratio. Unless stated otherwise, the training set is further split into train and validation sets, through a 5-fold cross-validation procedure (the splitting is again 80:20) (Arlot and Celisse, 2010). After we found the best model hyperparameters through an exhaustive grid search during CV, we retrain the model on the whole train + validation set. The final performances are calculated on the test set. Except when explicitly stated, we do not use methods to mitigate class imbalance.

We measure the performances in terms of accuracy (a), recall (r , often called sensitivity), specificity (s), precision (p) and F1 score (F_1). The definition of these indicators is provided in C.

2.5. Error estimates

When possible, in order to get an intuition on whether two results are significantly different, we compute error bars on the performances as the standard error between the cross-validation folds. These error bars represent the performance fluctuations of models trained only on the validation folds. Since the validation folds are smaller than the test set, the fluctuations on the test set should be smaller. In addition, while the performances on the validation folds are computed on models trained on the training folds, the performance on the test set comes from a model trained on training + validation data, so we can expect the test performance to fluctuate less or equally. While not rigorous, this procedure allows for an easier interpretation of the results. The errors bars are the number given in parenthesis and are applied to the last significant digit (e.g. 0.03(1) stands for 0.03 ± 0.01).

Table 1

Recall, specificity and balanced accuracy (average of recall and specificity) of repeated experiments, using two kinds of estimators defined in Section 2.6.

Metric	Classification	Recall	Specificity	Bal. Acc.	Acc
Upper Bound	Binary	0.950	0.963	0.956	0.958
Conditional	Binary	0.797	0.871	0.834	0.841
Upper Bound	Multi	0.833	-	-	0.878
Conditional	Multi	0.623	-	-	0.612

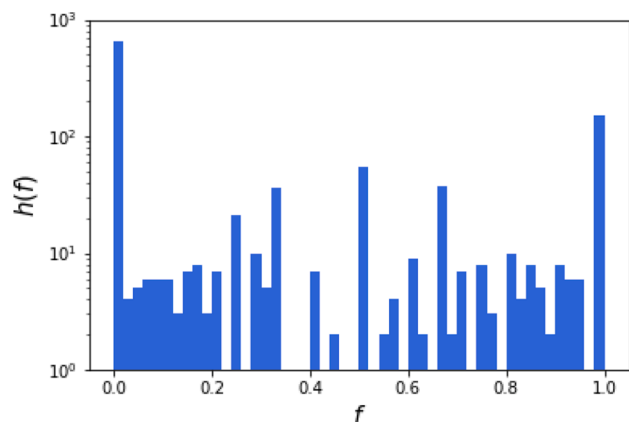


Fig. 2. Histogram $h(f)$. For each chemical, the quantity f is the fraction of species for which the given chemical is labeled as more toxic ($y_2 = 1$, after taking the median on the experiment repetitions as described in Section 2.2). If $f = 0, 1$, the chemical has the same binary outcome on all species, while if $f \neq 0, 1$, the chemical has a different impact according to the species.

2.6. Reproducibility of the experiments

Even though from the perspective of a regulator each experiment has one true outcome (in the binary case, more vs less toxic), if an experiment is performed more than once it sometimes happens that different repetitions of the experiment result in different values of the label y_2 (or y_5 , if we are in the multiclass setting). We can interpret this variability as some source of fluctuations – which derive from elements which are not completely under control – around a *ground truth* value.³

Since these fluctuations are present in our whole data set, also the test set will have some degree of disagreement with the ground truth. As a consequence, testing a model that perfectly predicts the ground truth will result in an accuracy \tilde{a} lower than 100%. This means that the maximum accuracy that we should expect for our models is not 100%, but \tilde{a} , so the performances that we find should be measured in units of \tilde{a} . However, we do not have access to \tilde{a} , so the best we can do is to try to estimate it by analyzing the mutual (dis) agreement of experiments that were performed more than once. We use two metrics: (a) an estimated upper bound to \tilde{a} , and (b) a conditional estimator proposed in Ref. (Luechtefeld et al., 2018).

2.6.1. Upper Bound

Intuitive definition. We estimate an upper bound, A_{up} , to the maximum achievable accuracy \tilde{a} . When all the repetitions of an experiment agree, we can assume that they correspond to the ground truth. When they do not coincide, the accuracy is maximized by assuming that the value of y_2 (or y_5) corresponding to the ground truth is the one occurring most often. We then define A_{up} as the accuracy obtained by assuming that

³ An alternative would be to use a Bayesian approach, which describes the fluctuations as intrinsic, and would aim for example at reconstructing the full probability distribution of the LC50.

repetitions of an experiment measure the ground truth more often than not. If we restrict the measurement of A_{up} to the experiments with a majority of positive outcomes, we obtain an upper bound to the recall, R_{up} , whereas if we restrict it to the negative outcomes, we obtain an upper bound to the specificity, S_{up} . The balanced accuracy is the average between R_{up} and S_{up} . In D we provide the formal definitions, for the binary and multiclass cases.

An example. To provide some intuition, let us focus on a single experiment ($n_{rep} = 1$), with binary labels, and assume that it is performed five times, with results $y_2 \equiv (y_{2,1}, y_{2,2}, y_{2,3}, y_{2,4}, y_{2,5}) = (0, 1, 0, 1, 1)$. Since the value $y_2 = 1$ appears most often, we assume that it corresponds to the ground truth, and obtain $A_{up} = 0.6$ (similarly, we could obtain a lower bound $A_{down} = 0.4$). Notice that, by construction, A_{up} cannot be smaller than 0.5.

2.6.2. Conditional Estimator

Definition. The conditional estimators of recall and specificity used in Ref. (Luechtefeld et al., 2018) aim at calculating the actual value of those quantities, instead of an upper bound. Their procedure postulates that one of the repetitions of the experiment is the true result, and validates the other repetitions assuming it as truth. For a formal definition we refer the reader to Ref. (Luechtefeld et al., 2018).

An example. We can use the previously defined example of an experiment repeated five times, with results $y_2 = (0, 1, 0, 1, 1)$. We take the first outcome, $y_{2,1} = 0$, as true, and see that the accuracy with the remaining experiments is $a_1 = 0.25$. We then iterate through all the outcomes, and obtain $\bar{a} = \frac{1}{5} \sum_i a_i = 0.4$. For this specific example, this number is equal to the lower bound A_{down} obtained in the previous section, so it seems like a pessimistic estimate of the consistency of repeated experiments. To estimate the recall r (or specificity s), we restrict to the examples with a positive (or negative) outcome. We therefore have an estimate $r = 0.5$ and $s = 0.25$.

Experiment reproducibility metrics in our dataset. In Table 1 we show the value of the upper bound and conditional estimator measured from our data.

Caveat. Note that these estimators \tilde{a} are measured through different repetitions of the same experiment, whereas our models are trained and tested on the medians of these repetitions, which arguably make the labeling more stable. Thus, it is possible that \tilde{a} is higher than what we would expect from our estimators.

3. Results

3.1. Dealing with many species

Compared with training and testing models on a single species, dealing with a large number of species requires a different perspective. On one side, [i] only a small number of species-chemical couples are present in this kind of dataset. In other words, an ecotoxicological dataset is sparse. If our dataset was dense, it would have $n_{spec} n_{chem} = 831222$ lines, instead of only 20128. On the other side, [ii] the different chemicals have different degrees of toxicity depending on which species is exposed to them.

We can test statement [ii] from our dataset. For each chemical that is tested on more than one species, we measure the fraction f of species for which it results more toxic (*i.e.* it has a label $y_2 = 1$). If a chemical has the same effect on all taxa, then $f = 0, 1$. We plot the histogram of $f, h(f)$, in Fig. 2, where it can be seen that a sizable number of chemicals have $0 < f < 1$, meaning that these chemicals act differently on different species. The two peaks at 0 and 1 are expected, since we still expect that a number of chemicals have a similar effect across all the orders (especially because our analysis is restricted to a narrow taxonomic range, *i.e.* fishes). Furthermore (i) many of the entries at $f = 0, 1$ come from chemicals which are tested on only a couple of species, and (ii) we did not filter f according to the taxonomic distance between species.

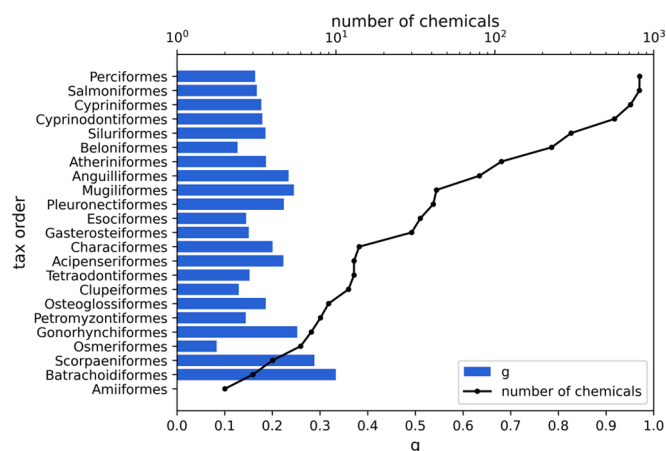


Fig. 3. The bars depict, for each order, the *taxonomic disagreement*, g , as defined in the main text. If $g = 0$ for an order, all chemicals tested on that order give the same result also on the other orders. If $g > 0$, that order reacts to chemicals differently from other orders. The black line (top x axis) says how many chemicals were tested on that order.

Table 2

Estimates using only the information on the chemical. Metrics of different models, trained and tested on \vec{x}_{ch} , i.e. without distinguishing among taxa. We highlight the best performance in bold. The numbers in parentheses are an uncertainty estimate (see Section 2.5).

Model	Accuracy	Recall	Specificity	F ₁
LR	0.841(5)	0.32(1)	0.969(6)	0.44(2)
3-NN	0.841(7)	0.46(2)	0.935(7)	0.53(3)
RF	0.855(5)	0.391(9)	0.969(8)	0.515(9)
MLP	0.84(2)	0.5(1)	0.93(1)	0.55(8)
S-RASAR	0.823(7)	0.32(2)	0.946(4)	0.42(2)
DF-RASAR	0.848(8)	0.30(3)	0.983(6)	0.44(3)

Table 3

Out-of-domain generalization, for different models, using only \vec{x}_{ch} . Here, models are trained on data coming from a single species (Rainbow Trout or Fathead Minnow), and they are tested on all the other taxa contained in the data set.

Model	Training Species	Acc.	Recall	Spec.	F ₁
1-NN	Fathead Minnow	0.805	0.668	0.880	0.709
RF	Fathead Minnow	0.814	0.647	0.906	0.712
MLP	Fathead Minnow	0.781	0.352	0.916	0.402
S-RASAR	Fathead Minnow	0.807	0.572	0.936	0.678
DF-RASAR	Fathead Minnow	0.817	0.606	0.933	0.702
1-NN	Rainbow Trout	0.820	0.827	0.816	0.753
RF	Rainbow Trout	0.843	0.800	0.865	0.773
MLP	Rainbow Trout	0.798	0.684	0.827	0.583
S-RASAR	Rainbow Trout	0.837	0.767	0.872	0.758
DF-RASAR	Rainbow Trout	0.830	0.788	0.850	0.755

Table 4

Estimates of y_2 from the whole dataset $\vec{x} = (\vec{x}_{ch}, \vec{x}_{tax}, \vec{x}_{ex})$, which combines information on chemical, taxonomy and experimental conditions.

Model	Accuracy	Recall	Specificity	F ₁
LR	0.882(2)	0.834(5)	0.910(2)	0.841(4)
3-NN	0.918(1)	0.888(5)	0.892(3)	0.890(2)
RF	0.932(1)	0.903(3)	0.949(3)	0.909(2)
MLP	0.913(4)	0.887(4)	0.929(6)	0.887(5)
S-RASAR	0.910(1)	0.858(7)	0.941(5)	0.877(2)
DF-RASAR	0.919(2)	0.861(4)	0.953(2)	0.888(3)

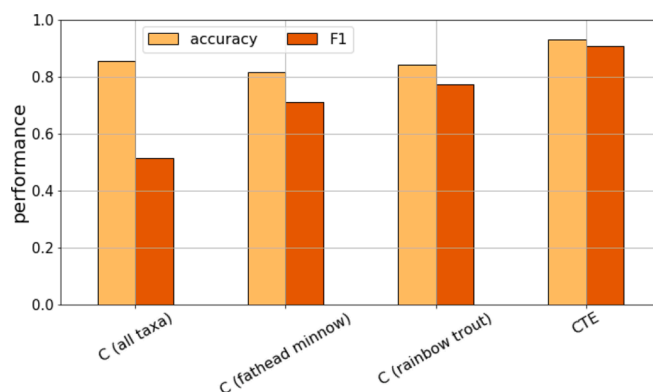


Fig. 4. Accuracy (red) and F1 score (green) of random forests trained with the \mathcal{C} setup (first three columns, data corresponding to Tables 2,3), and with the \mathcal{CTE} setup (fourth column, Table 4).

Table 5

Estimates of y_2 from the whole dataset $\vec{x} = (\vec{x}_{ch}, \vec{x}_{tax}, \vec{x}_{ex})$, using the distances defined in Eq. (5). We only list the models that rely on the definition of the distance [Eqs. (4) and (5)], since the others are unaffected.

Model	Accuracy	Recall	Specificity	F ₁
1-NN	0.918(2)	0.898(3)	0.931(4)	0.892(3)
3-NN	0.917(2)	0.896(2)	0.930(4)	0.890(2)
5-NN	0.908(4)	0.875(6)	0.928(3)	0.877(4)
S-RASAR	0.914(1)	0.896(6)	0.924(5)	0.886(1)
DF-RASAR	0.923(3)	0.884(5)	0.946(3)	0.908(6)

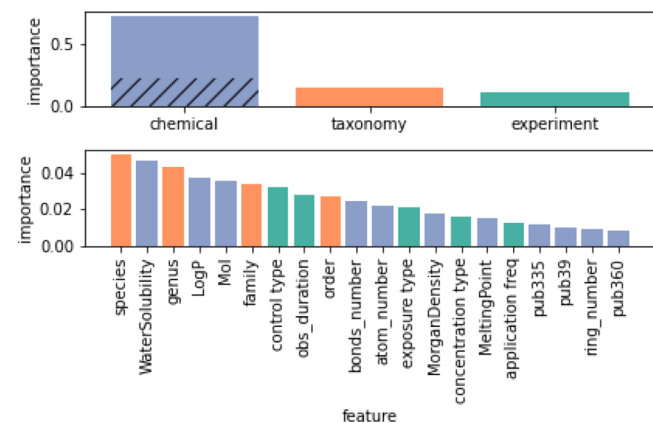


Fig. 5. Top: Feature importance from the RF. Every bar corresponds to a collection of features. The bottom (hatched) part of the bar corresponding to the chemicals is related to the chemical fingerprints, whereas the top part depicts the importance of the pubchems. **Bottom:** feature importance of the single features. For visualization purposes, we only include the 20 most important features (see A for a description).

The quantity f fixes the chemical, and checks how different the outcomes are on different species. We can try instead to fix the taxon, in order to see whether the outcomes on a specific taxon are very different from the others. To do this, for every taxon t , we take all the chemicals that were tested on t . Then, for each of these chemicals, we calculate the mean value of f , as defined in the previous paragraph. Finally, we calculate what we call the *taxonomic disagreement* parameter,

$$g = 1 - 2 \left| f - \frac{1}{2} \right|, \quad (6)$$

to indicate how different the outcomes on a taxon are from other taxa. If, for taxon t , $g = 0$, then all the chemicals that were tested on t had the

Table 6

Values of α_{Euc} , α_{pubc} and α_{cat} [Eq. (5)] for different models. These parameters are found during CV, and indicate the relative importance of the three sources of information (chemical fingerprints, pubchems, taxonomy & experimental details) for some of our models (see main text). For the RASAR models we use the ones of the 1-NN, since it is both the best K-NN model, and the one that is conceptually most similar to the RASAR.

Model	α_{Euc}	α_{pubc}	α_{cat}
1-NN	1	0.5623	0.0749
3-NN	1	1.7433	0.2212
5-NN	1	0.9459	0.0833
S-RASAR	1	0.5623	0.0749
DF-RASAR	1	0.5623	0.0749

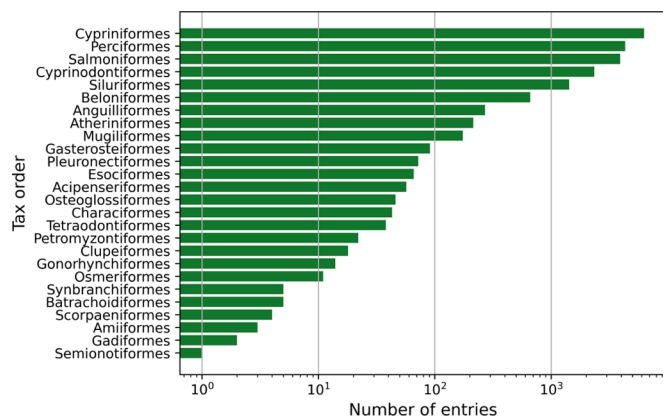


Fig. 6. Histogram of the different orders appearing in our dataset after data preparation.

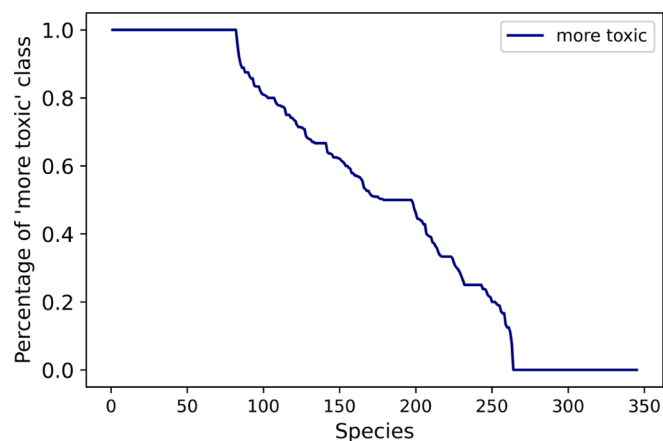


Fig. 7. For each species, we report the fraction of chemicals that results *more toxic* to it. Species on the x axis are ordered by rank, in order to have a non-increasing curve.

same outcome on all the other taxa on which they were tested. If $g > 0$, it means that part of the chemicals that were tested on t had a different outcome when tested on different values of the taxonomic rank. In other words, if it was useless to use taxa in our dataset, we would always have $g = 0$. The larger g (the maximum value is 1), the more crucial it is to use the taxonomic information.

In Fig. 3 we show the mean value of g for each taxonomic order (excluding those for which only a single experiment is available, see A). We chose the order for representation clarity, but as we explained this can be done in a similar way with any taxonomic rank.

The value of g is positive and steady across the orders, indicating that

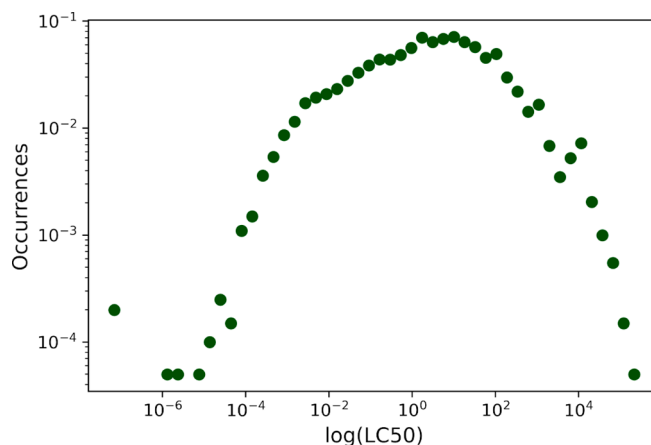


Fig. 8. Distribution of the LC50 values in our dataset.

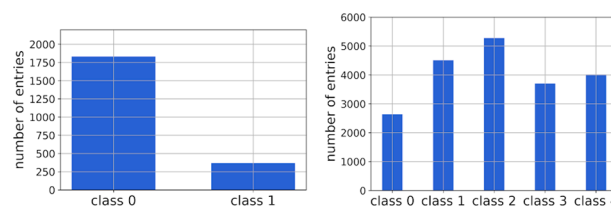


Fig. 9. Number of experiments per label after binary and 5-class labeling of the concentrations with only chemical features. The thresholds that define the classes are summarized in the appendix (Table 8).

Table 7

Optimal hyperparameters selected to construct the final MPL for each case as a result from hyperparameter tuning. Hyperparameters are: number of dense layers (Layers), number of nodes in each layer (Nodes), dropout rate (DR), and learning rate (LR).

Features	Classif.	Layers	Nodes	DR	LR
Only chemical	Binary	5	384,832,256,640,512	0.3	10^{-4}
Only chemical	Multi	4	448,1024,896	0.4	10^{-3}
Fathead Minn.	Binary	3	192,192,704	0.5	10^{-2}
Fathead Minn.	Multi	3	640,128,448	0.2	10^{-3}
Rainbow Trout	Binary	3	320,128,896	0.1	10^{-2}
Rainbow Trout	Multi	3	960,192,448	0.3	10^{-3}
All features	Binary	2	960,64	0.0	10^{-4}
All features	Multi	3	448,640,256	0.1	10^{-4}

experiments on one order are not fully representative of other orders, and that there is no single order that represents better or worse the other orders. The only exception would seem *Amiiformes*, but this datum is poorly significant because we only have two experiments on this order (the black line in Fig. 3 indicates the number of available experiments).

3.2. The \mathcal{E} setup: predictions based on chemical properties alone

Predicting toxicity outcomes based on chemical properties is the domain of quantitative structure–activity relationships (QSARs) (Cherkaev et al., 2014). Most often, existing QSARs are trained on the impact that a set of chemicals has on a single species. In our notation, this amounts to deducing y_2 from \vec{x}_{ch} . Approaches of this kind are sometimes called *global*, but since we refer to generic ways of using only \vec{x}_{ch} , we call this approach the \mathcal{E} setup. In this section, we show the outcome of this kind of approach on our dataset.

There are two ways in which we can train models using only \vec{x}_{ch} (and ignoring \vec{x}_{tax} and \vec{x}_{ex}):

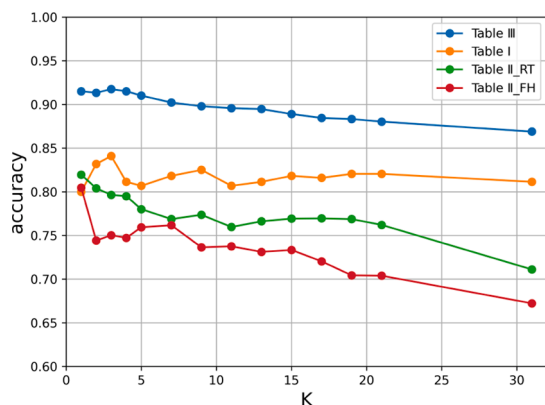


Fig. 10. Validation performance of the KNN models from Table 2 (using \vec{x}_{ch} without distinguishing among taxa), Table 3 (using \vec{x}_{ch} , training on a single species (Rainbow Trout or Fathead Minnow) and testing on all the other taxa) and Table 4 (ecotoxicological dataset \vec{x}) as a function of K.

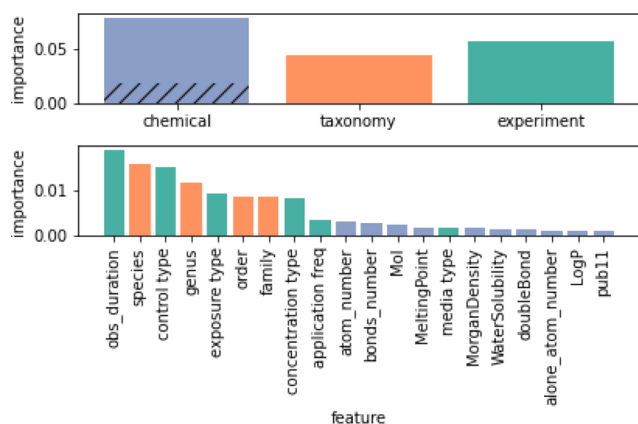


Fig. 11. Permutation-based feature importance from our best RF model. **Top:** Every bar corresponds to a collection of features. The bottom (hatched) part of the bar corresponding to the chemicals is related to the chemical fingerprints, whereas the top part depicts the importance of the pubchems. **Bottom:** feature importance of the single features. For visualization purposes, we only include the 20 most important features.

Table 8

LC50 intervals for the binary and multi-class labeling. For the binary labeling, we can say that the labels distinguish more (0) from less (1) toxic outcomes. For the multi-class, we can refer to the standard thresholds (Organisation for Economic Co-operation and Development, 2002; Li et al., 2017): non-toxic (0), slightly toxic (1), moderately toxic (2), highly toxic (3), very highly toxic (4).

LC50 interval (mg/l)	Binary Score	Multiclass Score
$(-\infty, 10^{-1}]$	1	4
$(10^{-1}, 10^0]$	1	3
$(10^0, 10^1]$	0	2
$(10^1, 10^2]$	0	1
$(10^2, +\infty)$	0	0

1. Train and test the models on the whole dataset, merging the \vec{x}_{ch} related to different taxa. The performances of the different models are shown in Table 2.
2. Train a model with data coming from a single species, and assess how well it generalizes to other species. As training taxa we use Rainbow Trout (*Oncorhynchus mykiss*) and Fathead Minnow (*Pimephales promelas*), which are commonly used (Ankley and Villeneuve, 2006)

Table 9

Metrics of different multi-class models, trained and tested on \vec{x}_{ch} , i.e. without distinguishing among taxa and experimental details (same as Table 2, but for multi-class classification.). The numbers in parentheses are an uncertainty estimate (see Section 2.5).

Model	Accuracy	Recall	Precision	F ₁
LR	0.45(1)	0.400(9)	0.45(3)	0.39(1)
1-NN	0.459(9)	0.459(5)	0.461(8)	0.458(6)
RF	0.545(9)	0.517(7)	0.55(1)	0.525(8)
MLP	0.47(3)	0.47(4)	0.48(2)	0.47(5)
S-RASAR(RF)	0.445(3)	0.422(7)	0.45(2)	0.426(8)
DF-RASAR	0.48(2)	0.46(2)	0.48(2)	0.46(2)

Table 10

Multi-class out-of-domain generalization, for different models, using only \vec{x}_{ch} (same as Table 3, but for multi-class classification). Here, models are trained on data coming from a single taxon (Rainbow Trout or Fathead Minnow), and they are tested on all the other taxa contained in the data set.

Model	Training Taxon	Acc.	Recall	Prec.	F ₁
1-NN	Fathead Minnow	0.493	0.497	0.503	0.498
RF	Fathead Minnow	0.512	0.518	0.525	0.518
MLP	Fathead Minnow	0.461	0.449	0.478	0.458
S-RASAR	Fathead Minnow	0.481	0.487	0.502	0.482
DF-RASAR	Fathead Minnow	0.498	0.505	0.515	0.501
1-NN	Rainbow Trout	0.523	0.532	0.532	0.525
RF	Rainbow Trout	0.541	0.545	0.553	0.541
MLP	Rainbow Trout	0.473	0.500	0.475	0.477
S-RASAR	Rainbow Trout	0.493	0.494	0.522	0.487
DF-RASAR	Rainbow Trout	0.523	0.526	0.539	0.522

Table 11

Performance of our models, trained on the full \mathcal{FE} feature space, in multi-class classification (same as Table 4, but for multi-class).

Model	Accuracy	Recall	Precision	F ₁
LR	0.639(5)	0.643(5)	0.641(6)	0.642(5)
1-NN	0.760(3)	0.765(3)	0.763(3)	0.764(3)
RF	0.780(2)	0.787(2)	0.784(2)	0.785(2)
MLP	0.78(1)	0.78(1)	0.78(1)	0.78(1)
S-RASAR	0.745(3)	0.750(2)	0.748(3)	0.749(2)
DF-RASAR	0.778(2)	0.784(2)	0.781(2)	0.782(2)

Table 12

Estimates of y_2 from the whole dataset $\vec{x} = (\vec{x}_{ch}, \vec{x}_{tax}, \vec{x}_{ex})$, using the distances defined in Eq. (5) (same as Table 5, but for multi-class).

Model	Accuracy	Recall	Precision	F ₁
1-NN	0.769(1)	0.775(1)	0.775(1)	0.774(1)
3-NN	0.751(4)	0.756(3)	0.753(4)	0.754(3)
5-NN	0.728(4)	0.729(4)	0.729(4)	0.729(4)
S-RASAR	0.770(2)	0.775(3)	0.773(2)	0.774(2)
DF-RASAR	0.781(3)	0.784(2)	0.783(3)	0.783(2)

Table 13

Values of α_{Euc} , α_{pubc} and α_{cat} [Eq. (5)] for different models for multiclass.

Model	α_{Euc}	α_{pubc}	α_{cat}
1-NN	1	0.3562	0.0127
3-NN	1	4.8939	0.0853
5-NN	1	3.0392	0.0672
S-RASAR	1	0.3562	0.0127
DF-RASAR	1	0.3562	0.0127

Table 14

Estimates of y_2 from the whole dataset with chemical stratification split using the \mathcal{C} approach, $\vec{x} = (\vec{x}_{\text{ch}})$.

Model	Accuracy	Recall	Specificity	F ₁
LR	0.810(8)	0.278(26)	0.946(6)	0.373(27)
2-NN	0.820(8)	0.209(13)	0.969(6)	0.313(16)
RF	0.818(7)	0.385(20)	0.931(7)	0.467(9)
S-RASAR	0.843(6)	0.261(7)	0.989(4)	0.400(9)
DF-RASAR	0.870(3)	0.338(16)	0.983(6)	0.477(14)

Table 15

Accuracy on binary classification with dataset stratified splitting by chemical stratification split using the \mathcal{CTE} approach, $\vec{x} = (\vec{x}_{\text{ch}}, \vec{x}_{\text{tax}}, \vec{x}_{\text{ex}})$.

Model	Accuracy	Recall	Specificity	F ₁
LR	0.789(23)	0.924(22)	0.590(47)	0.693(40)
3-NN	0.823(24)	0.751(35)	0.878(19)	0.785(42)
RF	0.764(21)	0.799(3)	0.699(48)	0.671(42)
S-RASAR	0.751(15)	0.867(11)	0.560(36)	0.629(39)
DF-RASAR	0.730(16)	0.948(24)	0.670(37)	0.671(26)

Table 16

Accuracy on binary classification with dataset stratified splitting by species stratification split using the \mathcal{CTE} approach, $\vec{x} = (\vec{x}_{\text{ch}}, \vec{x}_{\text{tax}}, \vec{x}_{\text{ex}})$.

Model	Accuracy	Recall	Specificity	F ₁
LR	0.852(10)	0.923(13)	0.678(10)	0.727(19)
9-NN	0.840(5)	0.774(24)	0.885(12)	0.797(18)
RF	0.895(10)	0.940(14)	0.834(21)	0.871(5)
S-RASAR	0.852(10)	0.927(10)	0.737(26)	0.797(19)
DF-RASAR	0.885(7)	0.964(12)	0.842(46)	0.861(39)

(we have 2500 and 3000 entries for each, see A). The performances of the different models are shown in Table 3.⁴

From Tables 2 and 3, it seems less advantageous to blindly mix all the taxa, than to train on a single species, and then naively extrapolate. The reason for this can be that the former method involves merging data coming from different taxa *before training*. This implies that i) the dataset becomes smaller and unbalanced (see A - anyhow, this can be partly offset by using the F1-score as a metric), and ii) that the labels y_2 are perturbed in an unknown way, since a label that correctly applies to one taxon does not necessarily apply to another. Further tests of the \mathcal{C} approach are shown in G.

3.3. The \mathcal{CTE} setup: combining chemical, taxon and experimental details

We now turn to what we define the \mathcal{CTE} setup, that combines all the different sources of information. At this stage, we treat \vec{x}_{ch} , \vec{x}_{tax} and \vec{x}_{ex} on the same footing. For KNN and RASAR we use the distance defined in Eq. (4) (details on model training can be found in B). In Table 4, we show the performance of each single model, and on Fig. 4 we plot the performance of RFs using the \mathcal{C} versus the \mathcal{CTE} approaches. The better performance of models trained with a \mathcal{CTE} approach stands out. Even a simple linear model such as the LR reaches good performances (F₁ \simeq 0.82) outperforming all the models obtained with both \mathcal{C} approaches (see Tables 2,3). Moreover, the gap between the LR and our best non-linear models is below 10%, demonstrating that much of the

⁴ We do not report error estimates in Table 3, because restricting the dataset to a single species made the training set significantly scarcer, so decided not to make a validation set and used the test set to select the hyperparameters. This means that the values in Table 3 are at risk of being overestimated, but it guarantees that they are not underestimated, which is our aim, as it will become clear in Section 3.3.

information can be captured through simple linear relationships. The best model is the RF, followed by the DF RASAR.⁵ The better performance, lower data requirement, and higher simplicity of implementation are strong factors in favor of the RF. However, we must keep in mind that the DF RASAR is explicitly designed for problems with a high availability of data from different biological effects, while in our dataset the effect “mortality” is dominant with respect to the other effects, which the DF RASAR uses as input to predict mortality (see A, and Appendix B). Furthermore, there is space for improving of RASAR models. One way is to replace the Euclidean distance (Eq. (4)) used by the RASAR with a distance that gives a different weight to different sources of information (Eq. (5)), as described in Section 2.3.2. In Table 5 we show the performances obtained by using this procedure. We note an improvement throughout most of the metrics of all the models (K-NN, S-RASAR and DF-RASAR), especially the DF-RASAR, which reaches $a = 92.3\%$ and F₁ = 90.8%.⁶

From Section 2.6 we know that the maximum achievable accuracy is not 100%, but rather an unknown value \tilde{a} . We cannot measure \tilde{a} directly, but we can estimate it by measuring A_{up} [Eq. (D.1)]. From Table 1, we see that the accuracy of the experiments is at most $A_{\text{up}} = 0.958$, so we can use this value to rescale our performances. For example, the rescaled performance of the RF becomes $0.932/0.958 \simeq 0.973 \approx 97\%$. On another side, we note that most of our models perform better than what would be indicated through the conditional estimator of \tilde{a} (Table 1), suggesting that this metric is a too pessimistic estimate of \tilde{a} .

3.4. Feature Importance

3.4.1. Importance from Random Forests

We now discuss the relative importance of the features, when using our best model (the RF of Table 4). These importances are estimates of how relevant each feature is for the predictions. Since there is no single best method to calculate feature importances, we evaluate them through two most commonly used methods: permutation-based and impurity-based (using Scikit-Learn libraries (Pedregosa et al., 2011)). While the permutation-based importance falls short with collinear features, the impurity-based tends to overestimate the relevance of high-cardinality categorical features (Altmann et al., 2010). Even though we do observe some (expected) differences between the two methods, the conclusions we draw are consistent across both kinds of analysis. We here show the results with the impurity-based method, and provide the permutation importance analysis in E.

In the top chart of Fig. 5, we coarsen the feature importances in three groups, corresponding to \vec{x}_{ch} , \vec{x}_{tax} and \vec{x}_{exp} . It can be clearly seen that \vec{x}_{tax} and \vec{x}_{exp} contribute to almost 30% of the information. The contribution of the features related to taxonomy and experimental details is even more striking when we look at them one by one: in the bottom set of Fig. 5, we show the importance of the single 20 most important features. The single most important feature appears to be the species. The most important features from \vec{x}_{ch} are water solubility and LogP. The control type and observation duration (details in A) appear to be the most influential experimental details, and the most important pubchems have the labels 335, which indicates the presence of the substructure C (~C) (~C) (~C) (~H), and 39, which indicates the

⁵ In G we provide further comparisons between the \mathcal{C} and \mathcal{CTE} approaches, showing that with different kinds of train/test splitting the \mathcal{CTE} approach still outperforms the \mathcal{C} , and that RF and DF-RASAR perform best.

⁶ Note that, even though, for completeness, we provide the KNN results for K = 1,3,5, comparisons between Tables 4 and 5 should compare KNN as a single algorithm. The specificity of S-RASAR shows a decrease but is compensated by the higher recall score, achieving a higher F1-score, which is generally the most desirable result.

presence of more than 4 C1 atoms ([Pubchem substructure fingerprint documentation](#)). Also with the permutation importance we find that the taxonomy- and experiment-related features, taken one by one, carry a large importance compared with the chemical descriptors taken one by one. These observations are consistent with Ref. ([Tuulaikhuu et al., 2017](#)), where the species is identified as the most important taxonomic descriptor, and the LogP as the most important chemical.

3.4.2. Importance from α

Since the distances used in Eq. (5) are normalized to one, the optimal values of each α tell us how relevant each set of features is. In particular, α_{cat} is an indicator of how relevant the data on experiment and taxon, \bar{x}_{ex} and \bar{x}_{tax} are with respect to the information on the chemical. In [Table 6](#) we show the optimal values of α for several models. The first thing that we point out is that α_{cat} is not zero, meaning that taxonomy and experimental details provide useful information. This is consistent with our previous feature importance analysis, and the remark that α_{cat} is ~ 20 times smaller than the other two descriptors suggests that these models rely more heavily on \bar{x}_{ch} than the RFs. Second, we notice that now the pubchems (\bar{x}_{pubc}) seem to often play a smaller role than the explicit chemical properties (\bar{x}_{pr}), than what we had seen with the RFs. We see two possible reasons for this: (i) the same information is contained both in the pubchems and in the properties, and which of the two is used is decided by the model while training. This is expected at least to a degree, since physicochemical properties can be predicted from the molecular structure ([Mansouri et al., 2018](#)). (ii) The distances in the pubchem space do not discern which chemical properties are different, and assign the same weight to all kinds of variation. This is also true for the chemical properties, \bar{x}_{pr} , but the properties are fewer, not binary, features. One way to come around this would be to treat different groups of pubchems components separately, by introducing more α constants.

3.5. Multiclass models

The whole analysis that we showed for binary classification was repeated in the case of 5-class labeling. Other than requiring to adjust the RASAR in order to also include an arbitrary number of classes ([B](#)), the analysis was analogous to the binary case. Also the results are similar, so we relegate them to [F](#). Here, we only make some remarks:

- The performances of our models are lower for multiclass, but also the self consistency of the data is lower, since with more classes it is more likely that the same experiment repeated twice falls into two different categories. If we normalize the recall by the upper bounds shown in [Table 1](#), we obtain a rescaled accuracy of 89%, and a rescaled recall of 95%.
- The performance of the multiclass DF-RASAR model matches that of the RF.

4. Conclusions

In this paper, we addressed the problem of predicting the acute toxicity of chemicals on different taxa through machine learning. We focused on fish datasets where one is interested in knowing about the effect of many chemicals on a large number of taxa, and the available data only includes few combinations of taxon and chemical. We addressed both binary and 5-class classification, with similar results.

Gain from including taxonomic and experiment information. We showed that taking into account the information on the taxon and on the experimental conditions is highly beneficial, providing over a 10% gain in F1 score (and other metrics, such as accuracy) with respect to the standard procedures of dealing with the same dataset. Remarkably, a linear model trained on a dataset which includes taxonomic information performs better than more sophisticated models which only make use of the chemical information. Additionally, we showed how one can take

into account the peculiarities of each source of data to obtain a further gain.

Best models and comparison to RASARs. The best performances are obtained by the RF, followed by the DF-RASAR and other models that obtained performances that were not too far. This similarity between different models was already found previously ([Li et al., 2017](#); [Hou et al., 2020](#)), and the superiority of RFs had already been pointed out ([Hou et al., 2020](#)). Even though the DF-RASAR obtains good performances, they do not seem to outperform other traditional machine learning models. In particular, the S-RASAR does not outperform K-NN, making its deployment less justifiable. The DF-RASAR, instead, has several advantages in the integration of the data from different effects, and a potential for improvement through some small variations.

Best performance in terms of maximum achievable performance. For binary classification, we obtain a maximum prediction accuracy of 93.2%. However, since this number comes from a test set which contains some inconsistent examples, this number is effectively higher than it appears, since it should be normalized by the accuracy of the test set. We cannot estimate it, but we see that it is roughly upper-bounded by $A_{\text{up}} = 0.958$. By rescaling appropriately through A_{up} , we can argue that the real accuracy is around 97%. Similar remarks are valid for the multi-class setting.

Metrics for comparing ML models with experiment reproducibility. If we use the metrics presented in Ref. ([Luechtefeld et al., 2018](#)), all our models (as long as we use a $\mathcal{L}_{\mathcal{L}}$ approach), outperform animal experiments. As also mentioned in Ref. ([Alves et al., 2019](#)), this statement should thus be polished, given the non-rigorous nature of this kind of estimators. To deal with this, instead of trying to estimate the accuracy (as well as other metrics) of the experiments, we estimated a heuristic upper bound. With these upper bounds, we still find evidence that our best models (RASAR, RF, KNN and MLP) are outperforming animal tests ([Section 2.6](#)). However, these comparisons should be taken with a grain of salt. In fact, the main problem is that, as stated earlier, this kind of comparisons is not rigorously defined. When we estimate the accuracy of the experiments, we have no truth value with which we can compare; we only have the results of repetitions of the same experiment. Therefore, we cannot estimate it without making strong assumptions. When we then calculate the accuracy of our models, we test them with a random selection of outcomes from experiments. In other words, we are now making the implicit (wrong) assumption that the data in the test set is not mislabeled (*i.e.* we are assuming that all the experimental outcomes faithfully reflect the ground truth). To overcome this issue, we made heuristic arguments on how the test accuracy of the models would be affected by the noise in the labels. However, these arguments are also not fully representative, since the accuracy of the labels is calculated on single repetitions of the same experiments (and assuming that all experiments fluctuate in the same way), but the data used for training and testing coarsens repeated entries in a single one, since we take the median LC50. Taking the median is desirable, since it arguably produces higher-quality data, but makes it even more far fetched to compare the performance of the models with that of the experiments. Anyhow, albeit not mathematically rigorous, these metrics arguably give a reasonable indication of the reliability of experimental results, and can be used to get a gross understanding of how accurate machine learning models are. In our case, they indicate that some of our models are close to the maximum obtainable performance, but only as long as we include taxonomic variability as an input variable. This means that, for the data we analyzed, the model selection is only marginally relevant with respect to the right choice and curation of the input features.

Most important descriptors. We discriminated which input features played the most important role. Most of the important features belong to the chemical, demonstrating why it is still possible to obtain decent predictions based on the chemical only. Nonetheless, the role of taxonomy and experimental details is relevant (*e.g.* it explains 30% of the inter-trees variability in the RFs), and the single most important feature is the species (the species descriptor is more crucial than the LogP

towards assessing toxicity), which is in agreement with previous observations (Tuulaikhuu et al., 2017). Moreover, in this work we restricted our analysis to fishes, but expanding the focus to wider taxonomic ranges would certainly further increase the importance of including the taxonomy.

Relation to human toxicology. Training models that are efficiently able to generalize across taxa is central in ecotoxicology, but it is also crucial for human-centered toxicology. In fact, the investigation of toxic effects on humans passes through the experimentation on mammals, such as rats, mice, guinea pigs and rabbits, and there is no guarantee that these species react in the same way as humans to novel toxicants. In recent years *in silico* methods, including machine learning, have been used to predict human toxicity - most often based on oral toxicity tested on rodents (Bhatarai et al., 2015; Gadaleta et al., 2019; Nedelcheva et al., 2019; Chushak et al., 2021; Mansouri et al., 2021), under the implicit assumption that humans react in the same way as the tested mammals. Therefore, mastering generalization across taxa can help data-driven approaches to infer better from animals to humans.

Replacing *in vivo* animal testing. In the long term, it would be desirable to replace *in vivo* testing through machine learning. How far are we from that? In order to do this, we would like *in silico* predictions to be as accurate, or almost, as directly performing an *in vivo* test. Our results represent a fraction of all the possible chemicals and taxa that one could be interested in testing. In order to expand the domain of applicability of our models, we would need to use a wider range of chemicals and taxa both for training and, more crucially, for testing. A known problem in machine learning is that the performances obtained in the test set often degrade when the models are deployed in real application settings, because the datasets used to test these models did not comprehensively represent the whole feature space. This problem is also known as dataset shift (Quiñonero-Candela et al., 2009; Moreno-Torres et al., 2012). One main source of dataset shift is the fact that the sampling of data points is not completely random, but follows patterns. This is certainly the case in (eco) toxicology, where taxa and chemicals that appear in the databases are chosen according to specific needs. Therefore, before the deployment of machine learning for risk assessment can become an option, it is necessary to have extensive-enough test sets, and a thorough study of dataset shift. Even then, it is foreseeable that for exotic chemical/

taxonomic categories the machine learning predictions might need to be confirmed through tests. Therefore, it is unlikely that machine learning alone will soon be able to completely replace animal testing, but it can certainly allow to reduce it drastically. In addition, given that it is anyhow hopeless to test all chemicals on all taxa, machine learning can be effectively used out of the box, even at the current stage of development, for prioritization: the results of *in silico* methods can guide us on which chemicals/taxa should be tested first.

Author Information

Jimeng Wu and Simone D'Ambrosi share first authorship.

Funding Sources

This project was funded by the SDSC grant "Enhancing Toxicological Testing through Machine Learning" (project N° C20-04).

Data and code availability statement

Our analyses were performed on an openly available data set (U.E.P. Agency, 2020). Our code for data cleaning and for model training and testing is freely available at github.com/mbaitye/ML-Tox.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

We thank C. Schür and L. Gasser for detailed feedback on the manuscript, M. Leone, G. Macchi and M. Vicentini for their preliminary results at the start of this project, and F. Balk, F. Hartig, M. Mair, F. Pérez-Cruz, P. Reichert, A. Scheidegger, N. Schuwirth and A. Zupanic for inspiring conversations.

Appendix A. Data

We use the Ecotox database, provided by the Environmental Protection Agency of the USA (U.E.P. Agency, 2020), and download the whole ASCII database. The dataset contains information on the taxonomy of the tested organism, on the conditions under which the experiment was brought through, a CAS identifier of the used chemical, and several endpoints related to the studied effects.

Chemicals. From the CAS codes provided in each EcoTox entry, we extract the SMILES string representing the chemical. We use the rdkit python library (RDKit) to extract the following features from the SMILES: number of atoms, number of alone atoms, number of OH groups, number of bonds, number of double bonds, number of triple bonds, and number of rings, LogP (octanol–water partition coefficient) and Morgan Density. Additionally, we extract molecular weight, water solubility and melting point from the Comptox database (Williams et al., 2017). From the SMILES we also obtain the Pubchem2D vector through the pubchempy library (Pubchempy). In a small number of cases we are not able to retrieve crucial information such as the SMILES or the Pubchem vectors. We decide to drop those data entries.

Taxonomy. We filter the dataset keeping only the experiments on fishes. This implies that we can discard the high-rank taxonomic features, ending up with 2 classes (Actinopterygii and Cephalaspidomorphi), 26 orders, 83 families, 219 genera and 345 species. We also exclude organisms at the embryo life stage. Note that the data is not equally distributed among taxonomic categories. For example, in Fig. 6 we show how many times every order appears in our dataset (after cleaning), which is highly imbalanced. We also show, in Fig. 7, the fraction of chemicals that result *more toxic* to it.

Experimental conditions. As for the experiment conditions, we keep those that have a sufficiently small number of missing entries: 'observation duration' (24 h, 48 h, 72 h, 96 h - most are 96 h), 'control type' (Concurrent control, Insufficient, Unknown, Satisfactory, Carrier or solvent, Multiple types, Unsatisfactory, Baseline, Positive controls), 'concentration type' (Active ingredient, Dissolved, Formulation, Labile [free metal ion], Not applicable, Not coded, Not reported, Total, Unionized), 'exposure type' (Static, Flow-through, Aquatic not reported, Renewal, Lentic, Pulse, Intra-peritoneal, Subcutaneous, Spray), 'application frequency', 'media type' (fresh water, salt water). The names are self-explanatory, and we refer to the Ecotox documentation for a detailed description of the features (<https://cfpub.epa.gov/ecotox/help.cfm?sub=term-appendix>).

Target. We then filter on the desired endpoint and effect. For mortality experiments, the endpoints EC50 and LC50 are equivalent, so we keep both. The distribution of the LC50 in our dataset is shown in Fig. 8. For the other effects, that are needed by the DF RASAR as input, we take the EC50.

Imputation and aggregation. We drop the lines with one or more missing features, and aggregate the experiments with the same CAS, taxonomy, and experiment conditions onto a single record, whose LC50 (or equivalent target) is the median LC50 from those repeated experiments. We then use the

median LC50 to calculate the labels y_2 and y_5 . While for the mortality data, which constitutes our target, we use the thresholds described in Eq. (1) (or Eq. (F.1) for multiclass), for the other effects, which are only used as input features for the DF-RASAR models, we use the median (or 20th, 40th, 60th and 80th percentiles for multiclass) in order to obtain maximally well-separated input features.

Standardization. All the features are rescaled between 0 and 1, and with a prior logarithmic transformation when the quantity spans several orders of magnitude. The scripts we used for the data cleaning are freely available (We provide our code at github.com/mbaitye/ML-Tox).

We end up with 20128 entries in our mortality dataset, including $n_{\text{chem}} = 2199$ chemicals and $n_{\text{spec}} = 345$ species.

The other studied effects (only for the DF-RASAR, see B) are genetics (777 entries), enzyme (691 entries), biochemistry (684 entries), accumulation (608 entries), behavior (422 entries), physiology (242 entries), cells (139 entries), intoxication (98 entries), and multiple (52 entries).

Class balance.

As shown in Fig. 1, the data set is balanced both in the binary and in the 5-class classification. If we neglect all the features different from \bar{x}_{ch} , we need to merge into a single entry all the lines related to different taxa/conditions and the same chemical, which makes the dataset less balanced. We end up having 461 positives and 1738 negatives and 543 non toxic, 640 hazardous, 555 toxic, 302 very toxic and 159 extremely toxic for multiclass (Fig. 9). If we restrict to fathead minnows, we have 2830 entries. If we restrict to rainbow trouts, we have 2508 entries.

Appendix B. Models

In this section we explain, where necessary, the implementations of our machine learning models. All our models are trained by splitting the dataset (80:20) into a training and a hold-out test set. The training set is used for 5-fold CV, implying a further 80:20 splitting into training and validation sets. The test set is not used until the very end, to produce the performance metrics.

Except for the RASARs, the models we use are quite standard, so we will not get into the details of most models, but rather in hyperparameter tuning and choices.

B.1. Logistic Regression and Random Forest

For LR and RF we use the standard implementation from the Scikit Learn python library (Pedregosa et al., 2011; We provide our code at github.com/mbaitye/ML-Tox). We rescale all numerical features through min–max normalization and perform one-hot encoding (using the OneHotEncoder function from Scikit Learn (Pedregosa et al., 2011)) on all categorical features in the preprocessing step. The hyperparameters adjusted in logistic regression include (the range of values in parentheses): C (10^{-3} - 10^3), $l1$ ratio (0–1), and max iter (100–800). In the random forest model, the adjusted hyperparameters are n estimators (200–1000), maximum depth (10–20), minimum sample segmentation (2–10), and minimum sample leaf (1–32).

B.2. Multi-Layer Perceptron

As a first preprocessing step, we one-hot encode all the categorical features. Then, we apply a min–max scaler to the data so that the scaled data has a minimum of 0 and a maximum of 1. The lower bound of 0 is chosen to keep the sparsity of the categorical features after one-hot encoding.

We use a tensorflow (Abadi et al., 2015) implementation of MLP and the tuning of hyperparameters is performed on the training set using the Hyperband tuning algorithm (Li et al., 2018). The tuned hyperparameters include (value range in brackets): number of layers (2–5), number of nodes in each layer (64–1024), dropout rate (0–0.5), and learning rate (10^{-5} - 10^{-2}). One dropout layer is placed after the dense layers and before the output layer. The tanh activation function is used in all dense layers. For the output layer, we use the sigmoid and the softmax activation for the binary and the multiclass classification, respectively. The selected hyperparameters for each case are shown in Table 7.

B.3. K Nearest Neighbors

The KNN algorithm assigns to an unknown entry the target that appears most often among the K nearest-neighbor data points. If two values are equally present, the one which is most present in the training set is chosen. For this reason, when the dataset is unbalanced (e.g. Table 2) even values of K give a higher performance. We use the KNN functions from SciKit Learn (Pedregosa et al., 2011).

We select and report for the value of K that gives the highest accuracy value. Differently from what is found in Ref. (Hou et al., 2020), we have a better test performance with small values of K. In Fig. 10 we show the validation performance for the different values of K that we tested, for both the \mathcal{C} and the \mathcal{CFE} setup.

B.4. RASAR

The RASAR model, presented in Ref. (Luechtefeld et al., 2018), comes in two variants. The S-RASAR requires data on a single effect (mortality in our case), whereas the DF-RASAR integrates several effects.

B.4.1. Simple RASAR

The S-RASAR is composed of two steps. In the first step, a metadataset is created through a nearest-neighbor algorithm, while in the second step, the metadataset is trained through a supervised learning method. The first step does not require any hyperparameter tuning, while the second follows the protocol described in B.1.

For each entry i in the training set, we save the normalized distance \tilde{d}_+ from the nearest positive,⁷ and the one from the nearest negative outcome, \tilde{d}_- (we use the distances defined in Section 2.3). This results in a meta entry

⁷ In Ref. (Luechtefeld et al., 2018), the similarity s with the nearest positive and negative is used. Since the distance \tilde{d} is normalized to 1, it is equivalent to use on the other, because $s = 1 - d$.

$$m_i^{(S)} = [\tilde{d}_i^+, \tilde{d}_i^-] \quad (\text{B.1})$$

We now have a supervised learning problem, where we need to deduce the labels y_{2i} from m_i . We did this through LR or RF.⁸

Multiclass. In Ref. (Luechtefeld et al., 2018), this algorithm is presented only as a binary classification algorithm, but we can easily extend it to multi-class classification. It is enough to measure the distance to the nearest element of each class. The multi-class metadata point is then

$$m_i^{(S)} = [\tilde{d}_i^1, \tilde{d}_i^2, \dots, \tilde{d}_i^C], \quad (\text{B.2})$$

where C is the number of classes.

B.4.2. Data Fusion RASAR

The DF-RASAR extends the idea of the S-RASAR, by including different sources of information in the metadataset. Since the definition in the binary case is already provided in Ref. (Luechtefeld et al., 2018), we define it here in the case of C classes.

The first, unsupervised, step is performed for E extra effects, beyond the target one, which we denote with a 0 (in our study, mortality). For each extra effect e , we construct an extra metadataset. This metadataset contains $C+1$ elements. The first C elements are the distances to the nearest experiment of each label (distance to the nearest 0, to the nearest 1, ...), and the $(C+1)^{\text{th}}$ is the label of effect e . In other words, if chemical A was tested on taxon B for mortality, now we are including the impact of A on B based on another effect e . If this information is unavailable, we use an extra "unknown" label. In the multiclass problem, the introduction of an "unknown" label implies that the six labels are no longer ordered and need to be treated as categorical variables.

In formulas, for each extra effect e , the metadata point corresponding to the i^{th} experiment (i.e. a taxon-chemical couple) is

$$\hat{m}_i(e) = [y_C(e), \tilde{d}_i^1(e), \tilde{d}_i^2(e), \dots, \tilde{d}_i^C(e)], \quad (\text{B.3})$$

where $y_C(e)$ is the C -class label (see F) of the e^{th} effect.

When the features used for the e^{th} effect are different from those of the target effect. In those cases, we restrict the distance to the common features.

As a result, the i^{th} meta datapoint for a C -class DF-RASAR takes the form

$$m_i^{(\text{DF})} = (m_i^{(S)}, \hat{m}_i(1), \dots, \hat{m}_i(E)) = \quad (\text{B.4})$$

$$= \left[\left(\begin{array}{c} \tilde{d}_i^1(0) : \tilde{d}_i^C(0) \\ y_{C_i}(0) \end{array} \right) - \left(\begin{array}{c} \tilde{d}_i^1(1) : \tilde{d}_i^C(1) \\ y_{C_i}(1) \end{array} \right) \dots \left(\begin{array}{c} \tilde{d}_i^1(E) : \tilde{d}_i^C(E) \\ y_{C_i}(E) \end{array} \right) \right]. \quad (\text{B.5})$$

We then use RFs on this metadataset to relate it to the target, $y_{C_i}(0)$.

Appendix C. Performance metrics

Here, we explicitly define the performance metrics used in this manuscript: accuracy (a), recall (r , often called sensitivity), specificity (s), precision (p) and F1 score (F_1).

The accuracy is the most intuitive of these quantities, and measures the fraction of correct answers,

$$a = \frac{\text{\#correct guesses}}{\text{total}}. \quad (\text{C.1})$$

In terms of the number of true positives (TP), true negatives (TN), false positives (FP) and false negatives (FN), the remaining metrics are defined as

$$r = \frac{TP}{TP + FN}, \quad (\text{C.2})$$

$$s = \frac{TN}{TN + FP}, \quad (\text{C.3})$$

$$p = \frac{TP}{TP + FP}, \quad (\text{C.4})$$

$$F_1 = 2 \frac{rp}{r + p}. \quad (\text{C.5})$$

⁸ Since we did not remark a difference between LR and RF, we only report the results using RFs, in order to be using the same supervised algorithm both in the S-RASAR and in the DF-RASAR.

The recall measures, out of all the positives in the test set, how many of those were guessed correctly. The specificity does the same, but for the negatives. In balanced datasets (*i.e.* with a similar number of positives and negatives), these are usually similar. With unbalanced datasets, high accuracies can be an artifact of the unbalance, and this usually results in very different values of r and s . In those cases, a better indicator of the performance is either the balanced accuracy ($b = \frac{r+s}{2}$), or the F_1 score (Luque et al., 2019).

For multiclass classification, we calculate the metrics on each class, and then average between the classes without taking into account how many examples there are in each class (this is called a macro average (Opitz and Burst, 2019)). In the multiclass case, the sensitivity does not make too much sense, since the negatives of one class include several classes. Therefore, we report the precision, which indicates, out of all the examples that were labeled as positive (for a given class) by the classifier, how many of those were actually correct.

Appendix D. Reproducibility of the experiments: formal definitions

Here, we define formally the quantities A_{up} , R_{up} , S_{up} and BA_{up} , that we introduced in Section 2.6. Let $y(i) = 0, \dots, C-1$ be a generic label (in this work we have C set to 2 or 5 classes), related to an experiment i , which is performed n_i times. Let us call n_{rep} the number of experiments that were performed at least twice, and $y_j(i)$ the label resulting from the j^{th} experiment. Let us also call $\bar{y}(i)$ the label of experiment i that appears the most times, over the n_i times that the experiment is performed. Then, we have

$$A_{up} \equiv \frac{1}{n_{rep}} \sum_{i=1}^{n_{rep}} \left[\frac{\sum_{j=1}^{n_i} \delta(y_j(i), \bar{y}(i))}{n_i} \right], \quad (D.1)$$

where $\delta(a, b)$ is the Kronecker function, which is equal to 1 if $a = b$, and equal to 0 if $a \neq b$. Therefore, the term in square brackets is the fraction of evaluations of the experiment that resulted in the most common outcome.

By restricting the measurement of A_{up} to the experiments with a majority of positive outcomes, we get an upper bound to the recall, R_{up} . By restricting to the negative outcomes, we obtain an upper bound to the specificity, S_{up} :

$$R_{up} \equiv \frac{\sum_{i=1}^{n_{rep}} \left[\frac{\delta(1, \bar{y}(i))}{n_i} \sum_{j=1}^{n_i} \delta(y_j(i), \bar{y}(i)) \right]}{\sum_{i=1}^{n_{rep}} \delta(1, \bar{y}(i))}, \quad (D.2)$$

$$S_{up} \equiv \frac{\sum_{i=1}^{n_{rep}} \left[\frac{\delta(0, \bar{y}(i))}{n_i} \sum_{j=1}^{n_i} \delta(y_j(i), \bar{y}(i)) \right]}{\sum_{i=1}^{n_{rep}} \delta(0, \bar{y}(i))}. \quad (D.3)$$

Here, are separating the experiments in two subsets. In Eq. (D.2) we are only considering the experiments with a majority of $y_j(i) = 1$ outcomes. This is enforced by the term $\delta(1, \bar{y}(i))$, which is zero unless the label is a 1. Since we discarded a fraction of the experiments, the normalization factor is no longer n_{rep} , but the total number of experiments with a majority of $y_j(i) = 1$ outcomes. The definition in Eq. (D.3) is equivalent, but we consider the experiments with a majority of $y_j(i) = 0$.

We obtain an estimated upper bound for the balanced accuracy by taking

$$BA_{up} \equiv \frac{R_{up} + S_{up}}{2}. \quad (D.4)$$

For multiclass classification, A_{up} is calculated in the same manner. Since in this case the specificity is not informative, we restrict to the recall. We calculate separately the average recall $R_{up,c}$ related to each class c , and then average over all the classes (*i.e.* we measure the macro-averaged recall):

$$R_{up}^{(MC)} = \frac{1}{C} \sum_{c=0}^{C-1} R_{up,c}. \quad (D.5)$$

Appendix E. Permutation-based feature importance

Since the impurity-based feature importance analysis is at risk of giving too much weight to high-cardinality features (Altmann et al., 2010), we also used permutation importance, to ensure that the species (which indeed has a high cardinality), indeed is one of the most important features. In the bottom set of Fig. 11 we see that, also when using permutation importance, the species remains one of the most important features. We notice that the importance of chemical-related features is now diminished with respect to the impurity-based method (this is also visible from the top chart). This can be expected, since the features in \vec{x}_{ch} are not linearly independent (*i.e.* similar information can be carried by more than one feature).

Appendix F. Multi-class performances

In this section, we show the results of the main paper, but with a multi-class target,

$$y_5 = \begin{cases} 4 & \text{if } c \leq 0.1 \text{ mg}/\ell, \\ 3 & \text{if } 0.1 \text{ mg}/\ell < c \leq 1 \text{ mg}/\ell, \\ 2 & \text{if } 1 \text{ mg}/\ell < c \leq 10 \text{ mg}/\ell, \\ 1 & \text{if } 10 \text{ mg}/\ell < c \leq 100 \text{ mg}/\ell, \\ 0 & \text{if } 100 \text{ mg}/\ell < c. \end{cases} \quad (\text{F.1})$$

The binary and multi-class (F) thresholds are summarized in Table 8.

In Tables 9 and 10 we report, for multi-class classification, the same results that in the main text we reported for binary classification (Tables 2,3, where we report the precision instead of the specificity, since it is better defined for multi-class classification). We observe the same trend shown for binary classification. The F1-score increase obtained by including taxonomy and experiment-related information is of around 15%. The best accuracy that we are able to reach with multi-class classification is 0.781 (Table 12), which is close to the estimators of the reproducibility of experiments ($A_{\text{up}} = 0.878$, Table 1), indicating that likely our models are close to achieving the best possible performance in this dataset. If we rescale the maximum obtained accuracy by A_{up} , we obtain $0.781/0.878 \approx 0.890 \approx 89\%$, which is close to we obtained for binary classification. The recall is even closer to its estimate maximum R_{up} , as it is $0.787/0.833 \approx 0.945 \approx 95\%$. (see Table 11)

Even in the multi-class case, the best models are RF, MLP and DF-RASAR, as it was for binary classification. From the values of the α parameters in Table 5 we find the same indications, with α_{cat} being smaller, though not negligibly, than α_{Euc} and α_{pubc} , and the pubchems seemingly less important when a smaller number of neighbors is taken into account.

Appendix G. Stratified splitting

Here, we show how our models perform in the case of a different splitting of the dataset. We choose two ways of doing stratified splitting: by chemical and by taxon. In the first, each CV fold, as well as the hold-out test set, contain exclusively different chemicals. In the second, they contain different taxa.

To do this, we used the GroupShuffleSplit function from SciKit Learn (Pedregosa et al., 2011) to split the train and test dataset, ensuring the (chemical, taxon) groups in the test dataset are independent of those in the training dataset. In the 5-fold CV, we used the GroupKFold function. For the LR and RF we adopt mitigate class imbalance through class reweighting.

G.1. Splitting by chemical

In Table 14 we show the performance of the \mathcal{C} approach when stratifying the performances by chemical, and in Table 15 we show the performance of the \mathcal{CTE} approach in the same conditions. The performances of the \mathcal{CTE} approach are better than the \mathcal{C} approach also in this case.

G.2. Splitting by species

In Table 16 we show the performance of the CTE approach when stratifying the performances by species.

References

- [SETAC] Society of Environmental Toxicology and Chemistry, 2018. Technical Issue Paper: Environmental Risk Assessment of Chemicals. Pensacola (FL): SETAC. 5 pp.
- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X., 2015. TensorFlow: Large-scale machine learning on heterogeneous systems, software available from tensorflow.org (2015). <https://www.tensorflow.org/>.
- Ai, H., Wu, X., Zhang, L., Qi, M., Zhao, Y., Zhao, Q., Zhao, J., Liu, H., 2019. Qsar modelling study of the bioconcentration factor and toxicity of organic compounds to aquatic organisms using machine learning and ensemble methods. *Ecotoxicol. Environ. Saf.* 179, 71–78. <https://doi.org/10.1016/j.ecoenv.2019.04.035> <https://www.sciencedirect.com/science/article/pii/S0147651319304580>.
- Alpaydin, E., 2020. *Introduction to machine learning*. MIT press.
- Altmann, A., Tološi, L., Sander, O., Lengauer, T., 2010. Permutation importance: a corrected feature importance measure. *Bioinformatics* 26 (10) (2010) 1340–1347. [arXiv:https://academic.oup.com/bioinformatics/article-pdf/26/10/1340/16892402/btq134.pdf](https://academic.oup.com/bioinformatics/article-pdf/26/10/1340/16892402/btq134.pdf), doi:10.1093/bioinformatics/btq134. doi:10.1093/bioinformatics/btq134.
- Alves, V.M., Borba, J., Capuzzi, S.J., Muratov, E., Andrade, C.H., Rusyn, I., Tropsha, A., 2019. Oy vey! a comment on “machine learning of toxicological big data enables read-across structure activity relationships outperforming animal test reproducibility”. *Toxicol. Sci.* 167 (1), 3–4.
- Ankley, G.T., Villeneuve, D.L., 2006. The fathead minnow in aquatic toxicology: past, present and future. *Aquat. Toxicol.* (Amsterdam, Netherlands) 78 (1) (2006) 91–102. doi:10.1016/j.aquatox.2006.01.018.
- Arlot, S., Celisse, A., 2010. A survey of cross-validation procedures for model selection. *Statistics surveys* 4, 40–79.
- Basant, N., Gupta, S., Singh, K.P., 2016. Modeling the toxicity of chemical pesticides in multiple test species using local and global qstr approaches. *Toxicology research* 5 (1), 340–353. <https://doi.org/10.1039/c5tx00321k>.
- Bhatarai, B., Wilson, D.M., Bartels, M.J., Chaudhuri, S., Price, P.S., Carney, E.W., 2015. Acute Toxicity Prediction in Multiple Species by Leveraging Mechanistic ToxCast Mitochondrial Inhibition Data and Simulation of Oral Bioavailability. *Toxicol. Sci.* 147 (2), 386–396. <https://doi.org/10.1093/toxsci/kfv135>.
- Cherkasov, A., Muratov, E.N., Fourches, D., Varnek, A., Baskin, I.I., Cronin, M., Dearden, J., Gramatica, P., Martin, Y.C., Todeschini, R., et al., 2014. Qsar modeling: where have you been? where are you going to? *Journal of medicinal chemistry* 57 (12), 4977–5010.
- Chushak, Y., Gearhart, J.M., Ott, D., 2021. In Silico Assessment of Acute Oral Toxicity for Mixtures. *Chem. Res. Toxicol.*, 34 (2), pp. 345–354. doi:10.1021/acs.chemrestox.0c00256.
- Furuhama, A., Hasunuma, K., Aoki, Y., 2015. Interspecies quantitative structure–activity–activity relationships (qsars) for prediction of acute aquatic toxicity of aromatic amines and phenols. *SAR QSAR Environ. Res.* 26 (4), 301–323.
- Gadaleta, D., Vuković, K., Toma, C., Lavado, G.J., Karmaus, A.L., Mansouri, K., Kleinstreuer, N.C., Benfenati, E., Roncaglioni, A., 2019. SAR and QSAR modeling of a large collection of LD50 rat acute oral toxicity data. *Journal of Cheminformatics* 11 (1), 58. <https://doi.org/10.1186/s13321-019-0383-2>.
- Goodfellow, I., Bengio, Y., Courville, A., Bengio, Y., 2016. *Deep learning*, Vol. 1. MIT press Cambridge.
- Hamm, J., Sullivan, K., Clippinger, A.J., Strickland, J., Bell, S., Bhatarai, B., Blaauboer, B., Casey, W., Dorman, D., Forsby, A., Garcia-Reyero, N., Gehen, S., Graepel, R., Hotchkiss, J., Lowit, A., Matheson, J., Reaves, E., Scarano, L., Sprankle, C., Tunkel, J., Wilson, D., Xia, M., Zhu, H., Allen, D., 2017. Alternative approaches for identifying acute systemic toxicity: Moving from research to regulatory testing. *Toxicology in vitro: an international journal published in association with BIBRA* 41, 245–259. <https://doi.org/10.1016/j.tiv.2017.01.004>.
- Hou, P., Jolliet, O., Zhu, J., Xu, M., 2020. Estimate ecotoxicity characterization factors for chemicals in life cycle assessment using machine learning models. *Environ. Int.* 135, 105393. <https://doi.org/10.1016/j.envint.2019.105393> <https://www.sciencedirect.com/science/article/pii/S016041201914412>. <https://reachonline.eu/clp/en/annex-i-4-4.1-4.1.2.html>.
- Li, F., Fan, D., Wang, H., Yang, H., Li, W., Tang, Y., Liu, G., 2017. In silico prediction of pesticide aquatic toxicity with chemical category approaches. *Toxicol. Res.* 6 (6), 831–842. <https://doi.org/10.1039/c7tx00144d>.
- Li, L., Jamieson, K., Rostamizadeh, A., Talwalkar, A., 2018. Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization, Tech. rep. (2018). <http://jmlr.org/papers/v18/16-558.html>.

- Luechtefeld, T., Marsh, D., Rowlands, C., Hartung, T., 2018. Machine learning of toxicological big data enables read-across structure activity relationships (rasar) outperforming animal test reproducibility. *Toxicol. Sci.* 165 (1), 198–212.
- Lu, J., Peng, J., Wang, J., Shen, Q., Bi, Y., Gong, L., Zheng, M., Luo, X., Zhu, W., Jiang, H., et al., 2014. Estimation of acute oral toxicity in rat using local lazy learning. *Journal of cheminformatics* 6 (1), 1–11.
- Luque, A., Carrasco, A., Martín, A., de las Heras, A., 2019. The impact of class imbalance in classification performance metrics based on the binary confusion matrix. *Pattern Recogn.* 91, 216–231.
- Mansouri, K., Grulke, C.M., Judson, R.S., Williams, A.J., 2018. Opera models for predicting physicochemical properties and environmental fate endpoints. *Journal of cheminformatics* 10 (1), 1–19.
- Mansouri, K., Karmaus, A.L., Fitzpatrick, J., Patlewicz, G., Pradeep, P., Alberga, D., Alepee, N., Allen, T.E.H., Allen, D., Alves, V.M., Andrade, C.H., Auernhammer, T.R., Ballabio, D., Bell, S., Benfenati, E., Bhattacharya, S., Bastos, J.V., Boyd, S., Brown, J. B., Capuzzi, S.J., Chushak, Y., Ciallella, H., Clark, A.M., Consonni, V., Daga, P.R., Ekins, S., Farag, S., Fedorov, M., Fourches, D., Gadaleta, D., Gao, F., Gearhart, J.M., Goh, G., Goodman, J.M., Grisoni, F., Grulke, C.M., Hartung, T., Hirn, M., Karpov, P., Korotcov, A., Lavado, G.J., Lawless, M., Li, X., Luechtefeld, T., Lunghini, F., Mangiardi, G.F., Marcou, G., Marsh, D., Martin, T., Mauri, A., Muratov, E.N., Myatt, G.J., Nguyen, D.-T., Nicolotti, O., Note, R., Pande, P., Parks, A.K., Peryea, T., Polash, A.H., Rallo, R., Roncaglioni, A., Rowlands, C., Ruiz, P., Russo, D.P., Sayed, A., Sayre, R., Sheils, T., Siegel, S., Silva, A.C., Simeonov, A., Sosnin, S., Southall, N., Strickland, J., Tang, Y., Teppen, B., Tetko, I.V., Thomas, D., Tkachenko, V., Todeschini, R., Toma, C., Tripodi, I., Trisciuzzi, D., Tropsha, A., Varnek, A., Vukovic, K., Wang, Z., Wang, L., Waters, K.M., Wedlake, A.J., Wijeyesakere, S.J., Wilson, D., Xiao, Z., Yang, H., Zahoransky-Kohalmi, G., Zakharov, A.V., Zhang, F.F., Zhang, Z., Zhao, T., Zhu, H., Zorn, K.M., Casey, W., Kleinstreuer, N.C., 2021. CATMoS: Collaborative Acute Toxicity Modeling Suite. *Environmental health perspectives* 129 (4) (2021) 47013. doi:10.1289/EHP8495.
- Miller, T.H., Gallidabino, M.D., MacRae, J.I., Hogstrand, C., Bury, N.R., Barron, L.P., Snape, J.R., Owen, S.F., 2018. Machine learning for environmental toxicology: a call for integration and innovation (2018). doi:10.1021/acs.est.8b05382.
- Moreno-Torres, J.G., Raeder, T., Alaiz-Rodríguez, R., Chawla, N.V., Herrera, F., 2012. A unifying view on dataset shift in classification. *Pattern Recogn.* 45 (1), 521–530. <https://doi.org/10.1016/j.patcog.2011.06.019> <https://www.sciencedirect.com/science/article/pii/S0031320311002901>.
- Nedelcheva, D., Stoeva, S., Dimitrov, S., Detroyer, A., Fadli, A., Note, R., Blanchet, D., Mekenyan, O., 2019. In silico mechanistically-based profiling module for acute oral toxicity. *Computational Toxicology* 12, 100109. <https://doi.org/10.1016/j.comtox.2019.100109>.
- Opitz, J., Burst, S., 2019. Macro fl and macro fl, arXiv preprint arXiv:1911.03347 (2019).
- Organisation for Economic Co-operation and Development, Guidance Document on the Use of the Harmonised System for the Classification of Chemicals which are Hazardous for the Aquatic Environment, OECD Publishing, 2002.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830.
- Pubchempy, <https://pypi.org/project/PubChemPy>.
- Pubchem substructure fingerprint documentation, ftp://ftp.ncbi.nlm.nih.gov/pubchem/specifications/pubchem_fingerprints.pdf.
- Quiñero-Candela, J., Sugiyama, M., Lawrence, N.D., Schwaighofer, A., 2009. *Dataset shift in machine learning*. MIT Press.
- RDKit: Open-source cheminformatics, <http://www.rdkit.org>.
- Rovida, C., Hartung, T., 2009. Re-evaluation of animal numbers and costs for in vivo tests to accomplish reach legislation requirements for chemicals—a report by the transatlantic think tank for toxicology (t4). *ALTEX-Alternatives to animal experimentation* 26 (3), 187–208.
- Sullivan, K., Allen, D.G., Clippinger, A.J., Wilson, D.M., Edwards, S.W., Glover, K., Mansouri, K., Settivari, R., Wijeyesakere, S.J., Casey, W., 2021. Mind the gaps: Prioritizing activities to meet regulatory needs for acute systemic lethality. <https://doi.org/10.14573/altex.2012121>.
- Tuulaikhuu, B.-A., Guasch, H., García-Berthou, E., 2017. Examining predictors of chemical toxicity in freshwater fish using the random forest technique. *Environ. Sci. Pollut. Res.* 24, 10172–10181. <https://doi.org/10.1007/s11356-017-8667-4>.
- U.E.P. Agency, Ecotox user guide: Ecotoxicology knowledgebase system. version 5.3, <http://www.epa.gov/ecotox/>, accessed: 2020-11-06.
- Vo, A.H., Van Vleet, T.R., Gupta, R.R., Liguori, M.J., Rao, M.S., 2019. An overview of machine learning and big data for drug toxicity evaluation. *Chemical research in toxicology* 33 (1), 20–37.
- We provide our code at github.com/mbaitye/ML-Tox.
- Williams, A.J., Grulke, C.M., Edwards, J., McEachran, A.D., Mansouri, K., Baker, N.C., Patlewicz, G., Shah, I., Wambaugh, J.F., Judson, R.S., et al., 2017. The comtox chemistry dashboard: a community data resource for environmental chemistry. *Journal of cheminformatics* 9 (1), 1–27.
- Wu, Y., Wang, G., 2018. Machine learning based toxicity prediction: from chemical structural description to transcriptome analysis. *International journal of molecular sciences* 19 (8), 2358.
- Xu, Y., Pei, J., Lai, L., 2017. Deep learning based regression and multiclass models for acute oral toxicity prediction with automatic chemical feature extraction. *J. Chem. Inform. Model.* 57 (11), 2672–2685.